

## LINK LAYER EMULATION

### Field of the Invention

The invention relates generally to communication networks. More particularly, the invention relates to communication link layer emulation.

### Background of the Invention

5       Packet networking is a form of data communication in which data packets are routed from a source device to a destination device. Packets can be networked directly between a source node and a destination node, or the packets can be relayed through a number of intermediate nodes. The packets generally include a source address, a destination address and a payload.

10       Figure 1 shows a network device 110 connected to a network 100. A client device 120 can also be connected to the network 100. For some networks, the network device 110 is a layer two (OSI model) device and the network 100 includes layer two routers. With the network 100 including only layer two devices, the network 100 appears to the network device 110 to be a transparent bridge. That is, if the network only includes layer two devices, then any client 120 connected to the network 100, appears to the network device 110, to be directly connected to the network device 110.

15

20       In some situations, it may be desirable to implement a network with layer three devices. Generally, a layer two device and layer two networks transparently bridge data packets without altering MAC or IP header fields of the data packets, whereas layer three devices generally alter MAC header fields of data packets passing through the layer three devices. The altering of MAC headers of data packets can present problems when interfacing a layer two device with a layer three network.

25       It is desirable to have a layer three network that allows a network device connected to the network to perceive the network as a transparent bridge. The network device should perceive a client connected to the network, to be connected directly to the network device.

**Summary of the Invention**

The invention includes an apparatus and method for emulating a layer two network. An embodiment of the invention includes a network of at least one router. The network includes a data packet address translator. The data packet address translator manipulates address information of data packets routed by the network, so that a network device connected to the network perceives the network to be a bridge. One embodiment of the network includes wireless mesh network that is wirelessly connectable to a client.

5 Other aspects and advantages of the present invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating 10 by way of example the principles of the invention.

**Brief Description of the Drawings**

Figure 1 shows a network device connected to a client through a network.

15 Figure 2 shows network device connected to clients through a network of routers, according to an embodiment of the invention.

Figure 3 shows network devices connected to clients through a network of routers, according to another embodiment of the invention.

Figure 4 is shows a network that includes a gateway and access nodes, according to an embodiment of the invention.

20 Figure 5 shows one example of a gateway according to an embodiment of the invention.

Figure 6 an upstream data path.

Figure 7 shows a downstream data path.

Figure 8 shows one example of a sequence of events for manipulating a MAC address of a downstream data packet.

Figure 9 shows one example of a sequence of events for manipulating a MAC address of an upstream data packet.

Figure 10 shows one example of a sequence of events for executing a pseudo proxy ARP.

### **Detailed Description**

5 As shown in the drawings for purposes of illustration, the invention is embodied in link layer emulation. More specifically, the invention can include emulating a layer two network, so that a network device connected to a layer three network, perceives that layer three network as a transparent bridge.

10 Figure 2 shows one example of a network device 210 connected to clients 220, 222, 224 through a network 230 of routers. The network 230 includes a MAC (media access control) address translator (MAT). The MAT of the network 230 provides manipulation of data packet addresses so that layer three routers of the network appear to the network device to be layer two devices. As a result, the network 230 appears to the network device 210 to be a transparent bridge. Additionally, the clients 220, 222, 224 appear to the network device 210 to be directly 15 connected to the network device 210. One embodiment of a layer 2 network is a transparent bridge.

The clients 220, 222, 224 can consist of a laptop computer, a personal digital assistant (PDA), a cell-phone, or any other device that includes an interface card adaptable for use with the networks of the invention.

20 The network device 210 can include any directly-upstream networking device (DUND). Exemplary DUNDs include routers, authentication gateways and switches. Generally, the network device 210 provides the client access to an internet or intranet network. The internet is a global network connecting millions of computers, in which control of the network is decentralized. Generally, each internet computer is independent, and referred to as a host. The 25 intranet is generally a network belonging to an organization that is only accessible by authorized

members of the organization. Generally, the intranet includes a firewall that surrounds the intranet to fend off unauthorized users.

Generally, client devices communicate with hosts on the internet or intranet. An upstream direction can be defined to be the direction away from a client device, whereas a downstream direction can be defined to be the direction towards a client device. A network device that is upstream of the network is defined as an upstream device. A network device that is downstream of the network is defined as a downstream device.

One example of a DUND is an authentication gateway. The authentication gateway provides authentication, authorization and accounting of client devices. Before a client can gain access to network resources provided by the DUND, the client must be authenticated by the authentication gateway. Many authentication protocols exist, including Radius, Diameter, ect.

Generally, the authentication gateway identifies the client devices based on the IP address of the client, or the MAC address of the client. Authentication based on the MAC address of a client can be problematic when interfacing a layer two DUND with a layer three network because the MAC address information of a client is not preserved when data packets of the client are routed through a layer three network.

The network 230 can be as simple as a single router, or the network 230 can include a mesh network that includes many routers. One example of a mesh network includes a wireless mesh network that can be wirelessly or wired connected to the clients 220, 222, 224. The wireless mesh network can include routers in the form of wireless access nodes. The wireless access nodes can provide wireless connectivity between access nodes, and to mobile clients.

Figure 3 shows one example of a first DUND 310 and a second DUND 312 that can be connected to the clients 220, 222, 224 through the network 230 of routers. A layer two switch 314 provides selectable connectivity between the DUNDS and the network 230. As shown in Figure 3, the first DUND 310 can be an authentication gateway, and the second DUND 312 can be a router.

Figure 4 shows another example of a network device 410 connected to clients 420, 422, 424 through a network 430 of routers. The network 430 of Figure 4 includes a gateway 440 and a number of wireless access nodes 450, 452, 454, 456. The wireless access nodes of Figure 4 operate as routers that are accessible by the clients 420, 422, 424 as shown, as well as other

5 clients not shown.

A gateway is a network entity that maintains an address mapping table for each client. As will be described, the address mapping table generally includes a MAC-IP address mapping for the client devices. A gateway can service several access nodes, or a gateway can be an access node. In this case, the gateway generally includes one or many more downlink interfaces.

10 An embodiment includes the gateway being an edge router between wired and wireless networks. An access node is a router that is directly connected (or connected through a layer two network) to one or more client devices.

The gateway 440 includes a MAC address translator (MAT) 441. The MAT 441 provides manipulation of the MAC addresses of data packets being routed by the network 430.

15 Link layer emulation by the gateway 440 (as provided by the MAT) provides that each packet sourced from a client device 420, 422, 424 and routed through the network 430 is received by the network device 410 with a MAC address that is the client devices MAC address. That is, the packet is received by the network device 410 with the client MAC address as the source MAC address, as happens when the network 430 includes layer two devices rather than layer three devices, or if the client device 420, 422, 424 is directly connected to the network device 410. Additionally, the network device 410 can use the actual destination MAC address of the receiving client device 420, 422, 424 when sending packets to the client device 420, 422, 424. The packets are received by the gateway 440 and forwarded through the layer three network to the correct client device 420, 422, 424. If the network device 410 issues an ARP (address resolution protocol) request to resolve a MAC address corresponding to the IP address of a client device, the gateway 440 responds on behalf of the client device with the correct MAC address of the client device.

#### Address Resolution Protocol (ARP)

The ARP protocol is a method for mapping an IP address to a corresponding MAC address. When a network device, such as a router receives a packet in which the destination IP address is not that of a local interface, the network device makes a forwarding decision. In the event that a direct route exists in the router's (network device) route table (indicating that the 5 destination IP address is directly attached to the router through one of its ports), the router attempts to forward the packet out over the interface on which the destination IP address is accessible (the downstream interface). Before transmitting the packet, however, the router creates and populates a MAC header for the packet. In order to do so, the router needs to know the destination MAC address corresponding to the destination IP address. Generally, the router 10 first consults a local ARP table that contains a mapping of IP addresses to MAC addresses. If a matching entry is found, the router populates the destination MAC address field of the packet with the value obtained from the ARP table. In the event that no entry is found corresponding to the destination IP address, the router issues a broadcast ARP request on the downstream interface for the destination IP address. The device whose IP address is this destination IP address 15 responds to the ARP request, providing device MAC address. With this information, the router is able to populate the MAC header field and forward the queued packet to the destination IP address. Information from the ARP response is used to construct an ARP table entry for the device Mac and IP address combination.

20

### Pseudo Proxy ARP

The DUND issues ARP requests for client devices. In this case, it is important for the gateway to respond on behalf of the client device with the correct MAC address of the client device. The client device cannot itself respond to the ARP request because the ARP requests are broadcast, and broadcasts are not forwarded by a layer three network. In order for the gateway 25 to respond with an ARP response containing the correct MAC address, the gateway needs to have a locally-maintained mapping of IP-MAC address for each client device. This is achieved through the operation of an AARP (Anti-ARP) protocol. The gateway consults an AARP database, determines the MAC address corresponding to the queried IP address and replies to the DUND with the correct MAC address. In order to achieve consistency, the source MAC address 30 of the ARP response is set to the MAC address of the client device whose MAC address is being

sought. This is important because certain DUNDS, do not update their ARP table when they receives an ARP response in which the source MAC address does not match the address in the body of the ARP frame. This process is performed by the gateway, and is referred to as a pseudo proxy ARP proxy (PPA) to distinguish this method from proxy ARP (in which the gateway 5 responds to the ARP request with its own MAC address).

In one embodiment, a software process listens on an open IP socket for ARP requests. When a request is received for an IP address that is identified as belonging in the local AARP database (thereby identifying a client device), the process constructs an ARP response with 10 source MAC address set to the MAC address of the client device, and the MAC address of the ARP response is set to the MAC address of the client device. This ARP response packet is transmitted on the same network interface on which it was received.

The DUND is able to associate the correct MAC address to the IP address of the client 15 device, as would happen if the network between the client device and the DUND had been a layer two network (in which broadcasts, such as ARP requests, are rebroadcast). This is one of the features of the link layer emulation mode.

#### Anti-ARP Protocol (AARP)

20 One or more gateways on the network can maintain a MAC-address-to-IP-address mapping for all client devices attached to the network. The address mapping can be stored in an address mapping table that can be arranged to be synchronized across all the gateways. This address mapping table can be referred to as an AARP Table (AARP = Anti-ARP). ARP is generally is a protocol for the resolution of IP addresses to MAC addresses. Here, anti-ARP 25 (AARP) is a protocol for the resolution of MAC addresses to IP addresses. The AARP protocol is based on a client-server architecture in which the AARP server on the gateway can be queried to extract the IP address corresponding to a given MAC address.

Generally, the network architecture includes a distributed network of AARP servers that may query each other and synchronize their address mapping tables periodically or on-demand in order to satisfy AARP queries.

5 Figure 5 shows one example of a gateway 440 with greater detail. The gateway 440 includes a MAT 510, a pseudo proxy ARP controller 520, an AARP controller 530, and routing logic 540.

10 The MAT 510 provides the MAC address manipulation required for the layer three network to appear as a layer two network to the DUND. More precisely, the MAT 510 modifies data packet MAC addresses to a client MAC address corresponding to an IP address of the data packets being routed by the network of the gateway 440.

15 The pseudo proxy ARP controller 520 provides a gateway response on behalf of a client device with the correct MAC address of the client device. In order for the gateway to respond with an ARP response containing the correct MAC address, the gateway needs to have a locally-maintained mapping of IP-MAC address for each client device. This is achieved through the operation of the AARP (Anti-ARP) protocol. The gateway consults an AARP database, determines the MAC address corresponding to the queried IP address and replies to the DUND 20 with the correct MAC address. In order to achieve consistency, the source MAC address of the ARP response is set to the MAC address of the client device whose MAC address is being sought.

25 The AARP controller 530 maintains a MAC-address-to-IP-address mapping for all client devices attached to the network. The address mapping can be stored in an address mapping table that can be arranged to be synchronized across all the gateways. The AARP is a protocol for the resolution of MAC addresses to IP addresses. The AARP protocol is based on a client-server architecture in which the AARP server on the gateway can be queried to extract the IP address corresponding to a given MAC address.

The routing logic 540 of the gateway provides the routing paths of data packets through the network, between the DUND and client devices. For a wireless access node mesh network, the route selection can be dynamic, and include several wireless hops between access nodes of the wireless mesh network. Generally, the creation, management, maintenance and deletion of 5 routes is accomplished through the operation of a routing protocol

### **MAC Address Translation (upstream traffic)**

A packet originating at a client device is forwarded through the layer three network using 10 IP forwarding. Typically management of routing tables occurs through the operation of protocols such as RIP or OSPF.

When the packet reaches the gateway, it is to be forwarded upstream to the DUND. Once the forwarding decision has been made, a MAC header should be constructed for the outgoing 15 packet (now on a queue). The destination MAC address is set to be the MAC address of the DUND, which is obtained either from the local ARP table or through the ARP request-response mechanism. The source MAC address would be, by default, be set to the MAC address of the outgoing interface on the gateway. This, however, is incorrect in link layer emulation mode. If the packet is sent with the MAC address of the gateway, the DUND receiving the packet would 20 update its ARP table, now associating the IP address of the client device with the MAC address of the gateway. This is not desired, and would result in service interruptions to the client device. Therefore the gateway should set the source MAC address to the MAC address of the client device. The gateway does this so (1) consulting an AARP database and extracting the MAC 25 address corresponding to the IP address in the source IP address field of the packet, (2) setting this MAC address as the source MAC address of the packet. The modified packet is now transmitted on the outgoing interface.

The net result of this manipulation of the packet is that the packet received by the DUND has source MAC address set to the MAC address of the client device and source IP address set to 30 the IP address of the client device – exactly as if the client device had been directly attached to

the DUND, or attached to the DUND through an intervening transparent layer two network. In one embodiment, a list of packet filtering rules is set up on the gateway to inspect each packet appearing on an incoming interface, detecting a match of the source IP address field with that of any of the IP addresses in the list, and modifying the forwarded (outgoing) packet so that the source MAC address field on the packet is altered to a MAC address specified by the list. The list of packet filtering rules in question is administered and set up by the AARP process. MAC address translation (downstream traffic).

Figure 6 shows an upstream data path for a data packet. For an upstream data path, the data packet originates at a client device, and is destined for a DUND 410. The data packet includes MAC source address and MAC destination address within a MAC header, IP source address and IP destination address within an IP header, and a payload. As the data packet travels from the client device, through an access node (for example AN 456), through a gateway (for example, gateway 440), to the DUND 410, the MAC addresses are manipulated by the network.

An originating data packet 620 includes a MAC source address that is the client's MAC address (Mclient), a MAC destination address that is the access node's MAC address (Man), an IP source address that is the client's IP address (IPclient), an IP destination address that is the destination's (DUND) IP address (IPdest), and a payload.

Upon being received by the access node 456, the MAC source address is modified to the access node's MAC address (Man), and the MAC destination address is modified to the gateway's MAC address (Mgw). This packet configuration is shown as packet 630.

Upon being received by the gateway 440, the MAC source address is modified to the client's MAC address (Mclient), and the MAC destination address is modified to the DUND's MAC address (Mdest). This packet configuration is shown as packet 640. Observation of the packet 640, reveals that the packet 640 appears to the DUND 410 to have been received directly from the client device.

### **MAC Address Translation (downstream traffic)**

5 When a DUND receives a packet for forwarding, it makes a forwarding decision and then constructs a MAC header using either its internal ARP table or through the ARP protocol. The resultant packet, if addressed to a client device, will have a destination MAC address that of the client device (this is ensured by virtue of the PPA mechanism discussed previously). This packet  
10 is received on the upstream interface of the gateway.

If a packet is received by a layer three device with destination MAC address not that of the layer three device or a broadcast or multicast address, the layer three device drops or discards the packet. In contrast, in one embodiment of the link layer emulation mode, the packet received  
15 by the gateway is first inspected to determine if the destination MAC address matches any of (1) the MAC address of the gateway, (2) the MAC address of a known client device. If a match is detected, the received packet is accepted for further processing. If the received packet has source MAC and IP addresses that match those of a client device (as determined from the AARP database), the packet is dispatched to the forwarding layer for a forwarding/routing decision. In  
20 another embodiment of link layer emulation mode, the received packet is intercepted by a packet filtering process that modifies the destination MAC address to the MAC address of the gateway. This has the effect that the drop/discard decision is never made and the packet always proceeds to the forwarding decision stage.

25 Figure 7 shows a downstream data path for a packet. For a downstream data path, the packet originates at the DUND 440, and is destined for a client. The data packet includes MAC source address and MAC destination address within a MAC header, IP source address and IP destination address within an IP header, and a payload. As the data packet travels from the DUND 410, through a gateway (for example, gateway 440), through an access node (for  
30 example AN 456), to the client device, the MAC addresses are manipulated by the network.

An originating data packet 720 includes a MAC source address that is the DUND's MAC address (Mdest), a MAC destination address that is the client's MAC address (Mclient), an IP source address that is the DUND's IP address (IPdest), an IP destination address that is the 5 client's IP address (IPclient), and a payload.

Upon being received by the gateway 440, the MAC source address is modified to the gateways's MAC address (Mgw), and the MAC destination address is modified to the access 10 nodes's MAC address (Man). This packet configuration is shown as packet 730. Observation of the packet 730, reveals that the packet 730 has been modified by the gateway so that the data packet will make it to the client without the DUND 440 being aware that the packet is passing through the network.

Upon being received by the access node 456, the MAC source address is modified to the 15 access node's MAC address (Man), and the MAC destination address is modified to the client's MAC address (Mclient). This packet configuration is shown as packet 740.

Figure 8 is a flow chart showing steps included within one example of a method of MAC address translation (MAT) for a downstream data packet. A first step 810 includes receiving a 20 packet with a destination IP address different from local interface addresses. A second step 820 includes making a decision. If the outgoing interface for the packet is not an upstream interface, a third step 830 is executed that includes sending the packet. If the outgoing interface for the packet is an upstream interface, a fourth step 840 is executed that includes determining if the source IP address matches an IP address within an AARP table. If there is not a match, a fifth 25 step 850 is executed that includes sending the packet. If there is a match, a sixth step 860 is executed that includes setting a source MAC address on the packet to the MAC address corresponding to the source IP address as determined from the AARP table. A seventh step 870 includes sending the packet.

Figure 9 is a flow chart showing steps included within one example of a method of MAC address translation (MAT) for an upstream data packet. A first step 910 includes receiving a packet on a gateway's upstream interface. A second step 920 includes determining whether the MAC address of the packet is an upstream interface MAC address. If the MAC address is an upstream interface MAC address, then a third step 930 is executed that includes accepting the packet for further processing. If the MAC address is not an upstream interface MAC address, then a fourth step 940 is executed that includes determining whether the destination MAC address of the packet is the same as a MAC address of a client device as determined from AARP tables. If the MAC address is not the same, then a fifth step 950 is executed that includes dropping the packet. If the MAC address is the same, then a sixth step 960 is executed that includes accepting the packet for further processing.

Figure 10 is a flow chart showing steps included within one example of a method of executing a pseudo proxy ARP. A first step 1010 includes an ARP request being received on a gateway's upstream interface. A second step 1020 includes determining whether the IP address of the request matches one of the gateway's interface addresses. If the request does match, then a third step 1030 is executed that includes sending an ARP response with the gateway interface MAC address. If the request doesn't match, then a fourth step 1040 is executed that includes determining whether the IP address of the ARP request matches an IP address of a client device as determined from an AARP table. If the request doesn't match, then a fifth step 1050 is executed including dropping the ARP request packet. If the request does match, then a sixth step 1060 is executed including constructing an ARP response with the MAC address set to the client device MAC address. A seventh step 1070 includes setting a source MAC address of the ARP response to be the client device MAC address. An eighth step 1080 includes sending the ARP response to DUND that sent the ARP request.

Although specific embodiments of the invention have been described and illustrated, the invention is not to be limited to the specific forms or arrangements of parts so described and illustrated. The invention is limited only by the appended claims.